World Development 121 (2019) 1–15

Contents lists available at ScienceDirect

World Development

journal homepage: www.elsevier.com/locate/worlddev

Quantifying household resilience with high frequency data: Temporal dynamics and methodological options

Erwin Knippenberg [a,*], Nathaniel Jensen [b], Mark Constas [c]

[a] Cooper/Smith, United States
[b] International Livestock Research Institute, Kenya
[c] Charles H. Dyson School of Applied Economics and Management, Cornell University, United States

ARTICLE INFO

Article history:
Accepted 17 April 2019

Keywords:
Resilience
Food security
Shocks
Machine learning
Africa
Malawi

ABSTRACT

Resilience as a metric is of growing interest to development researchers and practitioners, particularly for those whose work concerns the effects of climate change, conflict and epidemics. The growing need for resilience measurements motivates this research, measurements that reflect the complex, dynamic features of welfare among populations living in shock-prone contexts. It presents insights from three measurement approaches to explore the dynamic, intra-annual effects of shocks on household well-being, mediated by household characteristics; one based on shock persistence, one based on the stochastic distribution of well-being, and one driven by machine learning algorithms and based on predictive power. The insights gained from the comparison of measurement approaches offer a fuller understanding of the factors driving resilience than any single approach on its own. The paper harnesses a novel data-set from the 'Measuring Indicators for Resilience Analysis' project, which, each month, tracks shocks and food security indicators for one year in a highly food insecure population in Malawi. Across approaches, our study consistently finds that shocks and food insecurity are very persistent, and that households living in the flood plains are more resilient. When focusing on specific shocks such as illness, the gender of the household head matters as-well. A broader search for predictors of food insecurity using LASSO and Random Forest algorithms uncovers other characteristics that affect resilience, such as the distance to drinking-water. The paper demonstrates how multiple complementary approaches can identify different but overlapping factors related to resilience, painting a fuller, richer picture. It also demonstrates the empirical benefits derived from using high-frequency data sets in the study of resilience.

© 2019 Published by Elsevier Ltd.

1. Introduction

The world's poor face complex and interconnected risks. The poor are vulnerable to natural disasters driven by climate change, disease, conflict and global market forces, as-well as more idiosyncratic shocks such as death and divorce. Recognizing how this impedes the ability of households to climb out of poverty, policy makers and researchers increasingly emphasize the concept of resilience. Resilience as a concept is now widely used to formulate polices and interventions that can build long-term, sustained welfare gains. Investments are increasingly based on the idea that resilience can support the goal of helping households and communities better cope with an array of shocks and stressors undermining efforts to reduce poverty and food insecurity (Béné, Wood, Newsham, & Davies, 2012; Walsh-Dilley, Wolford, & McCarthy, 2016). The effort to improve resilience has led to a plethora of approaches to resilience measurement.

The importance of resilience as a key concept in contemporary academic and policy research is clearly illustrated in recent publications. For example, in the most recent four-year period, 2015 through 2018, 12 relevant articles featured resilience and resilience measurement in the pages of World Development (Angeon & Bates, 2015; Arouri, Nguyen, & Youssef, 2015; Bé né, Chowdhury, Rashid, Dhali, & Jahan, 2017; d'Errico & Giuseppe, 2018; Forsyth, 2018; Kuhl, 2018; Matin, Forrester, & Ensor, 2018; Quandt, 2018; Ramanath, 2016; Shah, Angeles, & Harris, 2017; Smith & Frankenberger, 2018; Thulstrup, 2015). In 2018 alone, World Development published six articles. With this body of work as a point of reference, resilience has emerged as a topic for development embraced by policy makers and by the research community. How then should researchers and policymakers make sense

* Corresponding author.
E-mail addresses: erwin@coopersmith.org (E. Knippenberg), N.Jensen@cgiar.org (N. Jensen), Mark.Constas@cornell.edu (M. Constas).

https://doi.org/10.1016/j.worlddev.2019.04.010
0305-750X/© 2019 Published by Elsevier Ltd.

of the many methods found in a rapidly growing literature and determine the resilience metrics that will best meet their needs?.

At the core of resilience measurement is a desire to unpack the relationship between well-being and shocks. For individuals living in shock-prone contexts, the dynamics that explains well-being are time dependent and connected to the experience of and recovery from shocks. This requires high-quality data collected at high frequency, by which we mean at least monthly. By leveraging a high-frequency data collection platform, the present paper explores how three different methodologies reveal different and complementary features of the relationship between shocks and well-being. These approaches are suited to different purposes, whether exploratory analysis, impact evaluation or targeting.

No papers that the authors are aware of calculate multiple resilience measures and compare the results. This research fills that gap, and it does so using a high-frequency panel of household data that was specifically developed and collected to track the impacts and persistence of shocks and linking them to welfare outcomes. Specifically, this paper identifies factors related to three conceptual approaches to measuring resilience drawn from the literature.

1. A household's assessment of the persistence of shocks on their own welfare. (Constas, Frankenberger, & Hoddinott, 2014a, 2014b).
2. The likelihood that a household's characteristics fall below a normative threshold of welfare. (Barrett & Constas, 2014; Cisse & Barrett, 2018).
3. The set of household characteristics that best predict future well-being. (McBride & Nichols, 2016; Kshirsagar, Wieczorek, Ramanathan, & Wells, 2017).

This paper uses a data set collected under the 'Measuring Indicators for Resilience Analysis' (MIRA) project. A team of researchers from Cornell and Catholic Relief Services (CRS) developed MIRA and piloted it in Malawi during a humanitarian emergency in 2016/2017. MIRA's main feature was a short, monthly, household survey using enumerators embedded in communities to track the adverse effects of self-reported shocks and indicators of food security. This high-frequency survey, which was preceded by a larger baseline survey, was administered each month for a year. The more detailed baseline survey allows us to investigate how time-invariant household characteristics affect the dynamics of shocks and food insecurity.[1]

This paper contributes to the literature on and understanding of resilience by presenting three approaches to its measurement in a single context drawing on aunified database, highlighting their complementary insights as well as their differences. First, this paper presents an approach emphasizing subjective shocks and their persistence to uncover relationships specific to particular types of shocks, such as the relationship between gender and resilience to illness. This is ideal for exploratory analyses intended to unpack household dynamics. Next, an approach presenting resilience as the probability of being above a threshold of well-being, in this case food security, allows for a model that accounts for multiple shocks and household characteristics. This approach is well-suited to impact evaluation. Finally, an approach focused on predicting food security as well-being using machine learning algorithms reveals potentially unexpected proxies, such as the distance to drinking water. By emphasizing predictive power, this approach is well suited to targeting exercises. The contrast among these methods is intended to illustrate that the resilience measurements overlap and complement each other. While each is tailored to a different purpose, all are focused on questions about how to model resilience dynamics. As an added methodological contribution, this paper highlights the need for monthly data on shocks and welfare in the resilience literature and describes how to collect such 'high-frequency' data centered on resilience.

The rest of the paper is organized as follows: Section 2 presents a brief review of the literature on resilience. Section 3 presents the data collection strategy and summary statistics. Section 4 outlines how transition probabilities can describe the persistence of subjective shocks' adverse effects. Section 5 compares the insights from this approach to the Cisse and Barrett (2018) moments based approach. Section 6 demonstrates the use of LASSO and Random Forest to infer the best predictors of future food insecurity. Section 7 discusses these three approaches, their complementary insights and recommended applications.

2. Literature review

The foundational literature on resilience spans many fields, including ecology, engineering and psychology. Holling (1973) characterized resilience as an ecological system's ability to remain or return to a dynamic equilibrium in the face of recurring shocks. Engineers view it as a physical system's ability to "mitigate hazards" (Tierney & Bruneau, 2007). Psychologists view it as 'adaptation to adversity' (Lee et al., 2013). A broad consensus exists that resilience focuses on the underlying parameters that determine the effect of a shock on an individual's current and future well-being.

In development economics, the emerging literature on resilience draws many of its core concepts from the poverty literature, which has shifted from positing poverty as a static feature to focusing on its dynamic and stochastic nature (Carter & Barrett, 2006). This includes acknowledging that poor households face high level of stochastic risk as their income and assets fluctuate, which has led to greater understanding of the vulnerability faced by poor households to weather and other shocks (Dercon, 2006). Faced with limited access to credit, insurance, and liquid assets, these vulnerable households often struggle to smooth consumption (Carter, Little, Mogues, & Negatu, 2007; Zimmerman & Carter, 2003), which leads to both transitory and long-term welfare losses as households are forced to forgo investments and sometimes reduce critical food intake (Barrett & Santos, 2014; Hoddinott & Kinsey, 2001). While the shift from static to dynamic representations of poverty represent a marked conceptual advancement, the events and conditions that trigger transitory and long-term losses have introduced new challenges for work on poverty. A deepened understanding of the pervasive effects of climate change (IPCC, 2018) across economic, socio-political and environmental domains requires a new approach for studying poverty dynamics of populations living in shock-prone contexts. Ideas contained within the concept of resilience are seen as providing this approach.

Arguably, the most comprehensive group of empirical studies on resilience measurement at the household level has been undertaken by the Food and Agriculture Organization (FAO) which began with studies of resilience in Palestine (Alinovi, Mane, & Romano, 2009). More recently, studies employing the FAO methodology have been conducted in Uganda and Tanzania (d'Errico, Romano, & Pietrelli, 2018), Senegal and Mauritania (d'Errico, Grazioli, & Pietrelli, 2018), and the Gaza Strip section of Palestine (Brück, d'Errico, & Pietrelli, 2018). Employing Principal Component Analysis and linear regression, Smith and Frankenberger (2018) have examined the shock-well-being dynamics in the contexts of floods in Bangladesh. While the work carried out by FAO, and the recent work of Smith and Frankenberger (2018) represents two important areas of work on resilience measurement, they do not examine

[1] 'Time-invariant' refers to slow-moving variables in the context of a single year.

comparative insights associated with different methodologies that might be used to model resilience.

Working with a panel of experts organized by the Food Security Information Network, Constas et al. (2014a, 2014b) present an early attempt to conceptualize how to measure resilience in the context of food insecurity. They emphasize the importance of focusing on dynamic well-being in the context of shocks conditional on characteristics, referred to as resilience capacities.This approach, which explicitly acknowledges that the effect of mitigating factors may depend on the shock in question, offers a causal model that emphasized the temporal dimension of resilience (Constas et al., 2014b). Illustrating the idea that resilience can also be characterized using time-series data, Smerlak and Vaitla (2016) take a purely non-parametric approach and consider a country resilient if its long-term food insecurity trend is non-negative and shocks experienced do not persist over time. In parallel Chavas (2017) uses a threshold quantile auto-regressive model and defines a resilient system as one where the unit root is greater than one during periods of low yield but drops below one when yields are at or above average.

In order to help articulate how the adverse effects of shocks stem from underlying uncertainty, Barrett and Constas (2014) argue that the concept of resilience provides a perspective on how to conceptualize and measure how and why stochastic well-being trajectories shift over time. In their work on development resilience, Barrett and Constas (2014) frame the concept of resilience as the changes in the probability distribution of well-being with respect to a normatively defined threshold, such as a poverty line or minimum consumption requirements. They define resilience as "the capacity over time of a person, household or other aggregate unit to avoid poverty in the face of various stressors and in the wake of myriad shocks. Resilience is presented conceptually as a probability drawn from the likelihood distribution of possible realizations of future well-being. In order to flesh this out empirically, Cisse and Barrett (2018) propose estimating the conditional mean and conditional variance of a household well-being indicator, such as assets, and constructing the conditional probability that this indicator will be above a given threshold. Their analysis then estimates the effects of household covariates and shocks on this probability, labelled 'resilience'.

Alternatively, one can think of resilience as a predictor of future well-being. Under this approach, resilient households are those less likely to experience low well-being over time as a function of both shocks and household characteristics. This emphasis on predictive proxies draws inspiration from an emerging interdisciplinary literature, which seeks to improve the accuracy of targeted social programs when lacking comprehensive data on income and consumption. This Proxy Means Testing (PMT) uses easily identifiable indicators, such as asset ownership, as proxies for poverty. While good at excluding the non-poor, such PMT-based formulas tend to miss many qualifying households. Brown, Ravallion, and Van De Walle (2018) show that a typical PMT-based formula applied to data from nine African countries can predict less than half of the extreme poor. But there have been attempts to improve targeting. For example, McBride and Nichols (2016) present evidence that applying machine learning algorithms to PMT development can substantially improve the out-of-sample performance of these targeting tools. Kshirsagar et al. (2017) similarly use a LASSO algorithm to identify and select the best predictors of poverty.

A recurring concern in the resilience literature is the lack of quality intra-annual data to track households' short-term experience of shocks and how this affects well-being. Data are most commonly collected as repeat cross-sections (e.g., census data) or, when longitudinal data are collected, they are usually at a resolution that is insufficient for observing the dynamics of shock and recovery in a world where shocks take place frequently and often in rapid succession or even simultaneously. Specifically, longitudinal 'high frequency' data enables researchers and policymakers to track household well-being and recovery from shocks at sufficiently high resolution so as to capture seasonality and short-term responses. A key constraint in collecting such data is the cost of repeatedly visiting households and the burden that frequent visits place on the households. To date, much of the existing high-frequency data stems from early warning systems, such as USAID's Famine Early Warning System Network, which, though useful in guiding humanitarian response, do not allow household-level analysis (Brown, Funk, Galu, & Choularton, 2007; Ververs, 2012). For these data to be useful for resilience analysis they would need to capture indicators on households' well-being, shocks, assets and demographic characteristics. Such data would also have to impose a sufficiently low burden for households to respond consistently on a monthly basis, imposing a maximum survey length of 10–15 min.

3. Data

3.1. The MIRA project

Malawi is a landlocked country in southeastern Africa. With fertile land and an influx of immigrants from its less stable neighbors, it has one of the highest population densities in the region. Eighty-four percent of its population lives in rural areas, and most of these rural people rely on subsistence agriculture (Bank, 2010). The combination of immigrants, high population densities, and agricultural livelihoods make Malawi particularly vulnerable to weather related shocks. As a case in point, the Shire river basin in southern Malawi was hit by devastating floods in January 2015, displacing hundreds of thousands of people. With resettlement underway, a consortium of development partners worked with the government to launch the United in Building and Advancing Life Expectations (UBALE) program, a program implemented by Catholic Relief Services (CRS) serving three of the poorest and disaster-prone districts in Malawi-Chikwawa, Nsanje, and Rural Blantyre.

In 2015, CRS and Cornell University agreed to collaborate on the development of a high frequency data collection protocol that would enable researchers and policymakers to track household food insecurity on a monthly basis at low-cost and with minimal burden on the participants. This differed from most 'early warning' systems in its panel structure, which permitted more sophisticated analysis than repeated cross-sectional data. In particular, this data set would permit the development of measures specific to shocks and coping capacities which CRS would use in its impact evaluation. With funding from CRS, the research team developed the Measuring Indicators for Resilience Analysis project.

The project was piloted in April 2016, commencing with a brief (45 min) baseline survey collected between May 18th and June 30th 2016. The baseline survey contained demographic, livelihood, economic, and shock history data. These household characteristics were considered either time invariant or sufficiently slow moving as to remain fixed over a year's time. In addition to standard indicators of assets, such as land and livestock, the survey included indicators of prosperity identified during pre-testing in the region. For example, a substantial minority of households who lived far from some of their fields had a secondary 'house', often a shed where they could store tools and sleep overnight if necessary. Data on the second house ownership therefore proved a useful proxy for spatial dispersion of the household's fields. We also included information on basic demographic indicators, including the incidence of chronic illness and disability.

The same households received monthly follow-up visits from June 2016 onwards, every month for a year up to and including May 2017. To facilitate the logistics of data collection, the team

trained local enumerators who lived in the communities and equipped them with smart-phones.[2] These embedded enumerators administered a 'high frequency' 5–15 min survey tracking the persistence of shocks and related food insecurity outcomes. This high-frequency data allows us to construct a panel, observing the dynamics of shocks and food security from month to month, linking them to underlying household characteristics. Importantly, the surveys retain respondents' prior information, allowing for follow-up questions that focused on the continued effects of previously reported shocks. This case management feature allowed us to more explicitly track the persistence of experienced shocks over time.[3] In order to minimize attrition, the researchers, in partnership with CRS, worked closely with community leaders to convey the importance of the data collection exercise. In addition, enumerators were instructed to be attentive and responsive to the preferences of respondents, scheduling visits in convenient locations and during periods that were not disruptive.

3.2. Sampling

The study takes place in the district of Chikwawa in sothern Malawi, where CRS is implementing the UBLA project. Within Chikwawa, we selected three traditional authorities (TAs)—Mikhwira, Ngabu, and Lundu— in the Shire river watershed to ensure variation in flooding history and risk. The Shire river flood plain, though more prone to flooding, is also more fertile than the higher lands surrounding it, and less prone to drought. In order to identify this flood plain objectively, we used flood-risk data from the Dartmouth Flood Observatory.

Within each TA, community level administrative units, called Group Village Heads (GVH), were stratified into high and low flood risk categories and then randomly sampled. This allowed for within TA, between GVH heterogeneity in how households experience drought and floods. During the roll-out an initially selected village was replaced with a village in an adjacent TA, M'bande, to maintain the balance between villages within the flood-plain and those outside the flood-plain. Fig. A1 in the supplementary materials illustrates the 2015 flood-zones and sampled households.

Two to three GHVs were randomly selected from each TA-strata for a total of 17 GHVs. One to two villages were then randomly selected from each GVH and 15–25 households were randomly selected from each village. The final sample was 576 households, from 31 villages, divided between high-risk flood-zones and low risk non flood-zones (Table 1).

With random selection carried out at the household level, the household is used as the unit of analysis.

3.3. Key Variables

The data used in this paper can be classified into three types: household characteristics X_i, shocks $Z_{i,t}$ and food insecurity $Y_{i,t}$. The baseline data include a series of household characteristics X_i which were not included in the high frequency survey. These include measures that other studies have found to be important and additional characteristics identified as important during preliminary field work. These include measures of assets, such as the amount of land farmed, measured in hectares, and livestock owned, measured using Tropical Livestock Units (TLU). Tropical Livestock Units are a composite index used to aggregate the number livestock a household owns, with weight equivalents for every species (Houérou et al., 1977). The following weights were used: cattle: 1, donkeys: 1, goats:.15, pigs:.2, chickens:.01.

Based on their geo-location and the extent of the 2015 flood, households were labeled as either living inside or outside the flood plain. Owning a second house was used as a proxy for having fields far from home, as members of the households would need to overnight there. Follow-up interviews confirmed that these secondary houses are used to sleep in while working in the further fields and usually little more than sheds, with negligible value as a standalone asset. The baseline also collected characteristics about the head of the household, including their age, gender and education level. It also recorded whether members of the members of the household reported themselves as chronically ill or disabled.[4] Table 2 summarizes these statistics. These are used to study how shock persistence and welfare trajectories differ across sub-groups.

The baseline questionnaire collected data on reported shocks $Z_{i,t}$, where shocks were defined as an unexpected event adversely affecting food insecurity.[5] Households were asked about a series of 15 shocks they may have experienced, as well as their perceived severity. The questionnaire prompted the enumerator to ask about previously reported shocks and whether households felt that they had recovered. If the household had not recovered, the questionnaire would prompt the same follow-up questions again in subsequent rounds, thereby creating a measurement of each shock's persistence as perceived by the injured household. Households were also asked about any new shocks experienced. A household could therefore experience the adverse effects of several shocks at once, and experience the same type of shock again after having previously recovered.

We chose to use subjective measures of shocks on the premise that households are best placed to report the impact of the shock on their own food insecurity. There is evidence to support that subjective measures track well with objective measures of well-being (Oswald & Wu, 2010; Stevenson & Wolfers, 2013). Furthermore, in cases where recovery is the outcome of interest, it is not clear how an outside researcher would weigh the various components that go into 'recovery', including human capital, assets, and social obligations. While communities are exposed to similar level of objective risk, their responses and perceptions are highly heterogeneous (Barrett, Smith, & Box, 2001; Doss, McPeak, & Barrett, 2008). Realized shocks are also heterogeneously perceived by the community, subject to individual reference points (Hunter, Gray, & Edwards, 2013). As it reflects a household's perception of a shock rather than its objective incidence, the measure is inherently endogenous to a household's capacity to cope and recover, a key feature of resilience. For example, a household may experience the effects of drought long after the objectively measured drought is over. In this context, observing the trends in the incidence and persistence of a shock's adverse effects can inform us about trends in households' well-being and coping capacity.

We collected data on 15 different shocks, including both idiosyncratic and co-variate shocks.[6] Fig. 1 shows the reported

[2] Tracking households month-on-month was facilitated by relying on local enumerators who knew the community. These enumerators were incentivized by a compensation scheme which paid them per successful survey uploaded, reducing attrition. To prevent fraud, the field supervisor in collaboration with the researchers monitored their meta-data, including location and time of survey.

[3] The open source Computer Assisted Personal Interviewing (CAPI) CommCare survey application was used for the high frequency survey because of its case-management functionality, which allows for a dynamic survey based on previous response. Surveys saved on enumerator's smart-phones were uploaded to the cloud every month, making household data available in near real time to researchers for prompt analysis.

[4] For confidentiality reasons we did not explicitly ask about HIV status. However conversations with enumerators and local health officials suggest that 'chronically ill' was often understood as implicitly referring to HIV.

[5] Potentially positive shocks were excluded as the effects on food insecurity are asymmetrical (Taylor, 1991).

[6] The total list of shocks accounted for are: drought, flood, crop disease, livestock disease, business failure, job loss, termination of assistance, fall in sales price, rise in food prices, household death, household breakup, illness, theft, wind damage and fire damage, as-well as an 'other' category. The four most frequent shocks were used in analysis.

Table 1
MIRA Study Sample.

Traditional Authority	GVH Strata (Flood Risk)	General Village Head (N = 17)	Villages (N = 31)	Households (N = 576)
Mikhwira	High	Mpama	2	37
	N = 103	Kanyimbiri	2	32
		Salvala	2	34
	Low	Nyambalo	2	38
	N = 98	Chagambatuka	2	37
		Champhanda	2	23
Ngabu	High	Jombo	2	47
	N = 82	Nkhwazi	2	35
	Low	Malikopo	2	38
	N = 112	Kalulu	2	44
		Chapomoko	2	30
Lundu	High	Mafale	2	37
	N = 105	Biliati	2	39
		Sekeni	1	29
	Low	Bestala	2	38
	N = 58	Biyasi	1	20
Maseya	Low	M'bande	1	21
	N = 18			
Totals	High risk	8	15	290
by Risk	Low risk	9	16	286

Table 2
Household Co-variates.

Characteristic	N	Mean	Std. Dev.	Min	Max
Land (Ha)	576	2.38	1.92	0	20
Tropical Livestock Units	576	0.57	2.42	0	38.04
Lives in Flood Plain (1 = Yes)	576	0.50	0.50	0	1
Secondary house or shed* (1 = Yes)	576	0.18	0.39	0	1
Head of Household:					
Age (Years)	576	42.14	16.22	0	97
Gender (1 = female)	576	0.25	0.43	0	1
Education (Years)	576	6.41	4.08	1	15
Chronically ill or disabled	576	0.17	0.38	0	1

* An indicator of owning fields far from home.

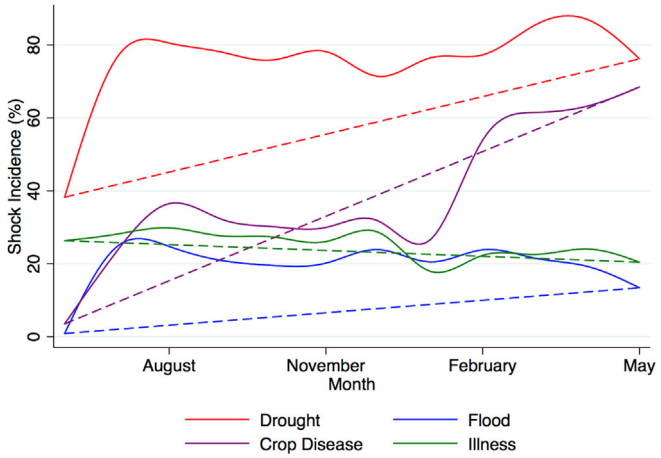

Fig. 1. Incidence of most frequent shocks reported, with the dotted line representing the trend lines if we were to extrapolate from the first to the last round.

household incidence of the most frequently reported subjective shocks across the 12 months monitored. These include drought, flooding, illness and crop-disease. The dotted line represents the trend lines if we were to extrapolate from the first to the last round and illustrates how much of the seasonal variation in shock experienced would go unobserved if we were to rely solely on data collected annually.

Food insecurity $Y_{i,t}$ is measured using the Coping Strategy Index (CSI). CSI is useful for rapidly measuring food insecurity in a humanitarian context, and is strongly correlated with more complex and time intensive measures of food insecurity (Maxwell, Caldwell, & Langworthy, 2008) It is a composite weighted score of various strategies households engage in when faced with short-term food shortages (Maxwell, 1996). Individual coping strategies c reflect activities households may be compelled into, often due to food insecurity, and compose the set $c \in C$. These include borrowing food, taking on piece work for additional income, consuming less preferred foods, reducing either the number or the size of meals and in extreme cases sending children to beg. The survey asks the number of days in the past week a household engaged in each of these activities, then multiplies those days by a weight w_c.[7]

CSI is therefore the weighted sum of days engaged in each coping strategy c:

$$CSI = \sum_{c}^{C} w_c * days_c \quad (1)$$

where $days_c$ is the number of days a household had engaged in a given coping strategy c over the past week, and w_c is the assigned severity weight.

[7] The following weights were used: Borrow food = 2, Piece Work = 1, consuming less preferred foods = 1, reducing meals = 1, reduce size of meals = 1, children begging = 4. These recommended weights are the result of extensive consultation and calibration (Maxwell et al., 2008).

A higher CSI score indicates higher food insecurity and therefore lower well-being. A household with a CSI of 10 may do some piece work on the side, eat less preferred foods or limit portion size a few days a week. A household with a CSI of 30 may do this every day, while also skipping meals and occasionally borrowing food. A household with a CSI of 60 is engaging in all these coping mechanisms daily, but must also send its children out to beg on occasion. A household engaged in all coping strategies all the time has a maximum CSI score of 70. In the context of chronic food insecurity, as in the Shire river basin, we consider CSI a valid measure of negative wellbeing.

For illustrative purposes, the observed trajectory of CSI was disaggregated according to household characteristics. For binary variables, the sample was disaggregated by type. For continuous variables, the sample was disaggregated by whether a household is above or below the median.

Figs. A2 and A3 in the supplementary materials illustrate the non-parametric CSI trajectories disaggregated by these observable characteristics. The data collection exercise began in the midst of a food emergency, with high levels of CSI throughout the population. The severity of the emergency, as measured using CSI, abated by the end of the year, but not equally for all groups. There is no appreciable difference in the CSI trajectory when dis-aggregated by the household head's age, years of education, or whether a member of the household is chronically ill. Households headed by men are worse off at first, but the trend quickly converges with households headed by women.

Households with above median Tropical Livestock Units (TLU) experience lower levels of CSI overall, but their trends mimic those of their neighbors who have below median TLU. There is a marginal difference for households with above median access to land, but the difference is not economically significant. Households living in the flood plain recover much faster and see their CSI drop. Having a secondary house, and therefore spatially dispersed fields, makes one marginally less food insecure, but the trends match those of households without a secondary house.

3.4. Attrition

Attrition is a recurring concern when collecting panel data. If too many participants drop out of the sample, the statistical power of estimates is weakened. We therefore over-sampled initially, allowing for up to 5% monthly attrition. Non-random attrition can also be a concern.This is problematic when it leads to correlation between the error term and the observables, potentially leading to biased estimates.

Table A1 in the supplementary materials illustrates attrition across time. The monthly attrition rate was 1.38% on average, though much higher in certain months. Between June and July logistical friction due to a change in enumerators and data collection platform meant a village was missed. In December, seasonal flooding washed out the roads, and an enumerator passed away. To mitigate attrition due to these events, households missed in a given round were still sought out for interviews in subsequent rounds, allowing them to re-enter the sample rather than drop out entirely.

To control for potentially non-random attrition, we use a Heckman style two-step estimator (Heckman, 1979). Observations are labeled as $Missing_{i,t} = 1$ if there are no observations for household i in period t, and $Missing_{i,t} = 0$ otherwise. Table A2 in the supplementary material reports the estimates from a probit maximum likelihood model to test whether this attrition was driven by observable characteristics. Since this is a panel, month fixed effects and community fixed effects are controlled for.[8] Finally, to address concerns regarding selection on unobservables, enumerator fixed effects are included in the selection equation as an exclusion restriction. The exclusion restriction is valid under the assumption that the identity of the enumerator affects the likelihood of response-variation in enumerator skill and personality will affect attrition. An F-test shows that the enumerator identity is a strong and significant predictor of whether an observation is missing, satisfying the inclusion criteria. The exclusion restriction is satisfied if enumerator identity does not systematically bias the reported shocks or well-being of households interviewed. Living in a flood plain is positively correlated with the probability of a missing observation. This discrepancy was anticipated when the sample was stratified. Other coefficients become insignificant once controlling for community and enumerator fixed effects. The predicted probability of selection is used to generate a Heckman inverse Mills ratio $\hat{\lambda}$, which is used as a control in subsequent specifications, similarly to Lillard and Panis (1998).

4. Resilience as perceived shock persistence

Building on the importance of the subjective shocks experience as argued above, we propose the perceived persistence of a shock's effects as a useful indicator of resilience. For example, imagine two households experiencing the same shock in a given month; in the next month, if one household is still experiencing the effects of the shock while the other has fully recovered, then the latter household is more resilient. All else being equal, the greater the shock's persistence, the lower the household's resilience to that particular shock.

With 12 rounds of data, we can track the persistence of individual subjective shocks over time. We posit two states $Z^s_{i,t} \in \{0, 1\}$, reflecting whether household i is experiencing the adverse effects of a subjective shock $s \in S$ in period t. Using the questionnaire's dynamic nature, respondents were prompted on the persistent effects of previously reported shocks if $Z^s_{i,t-1} = 1$. If households reported a full recovery, then $Z^s_{i,t} = 0$. If a household reported that they had not yet recovered from the given shock, then the shock's effects persisted and $Z^s_{i,t} = 1$. In the next round households were prompted about the effects of ongoing shocks, as well as whether they experienced new ones. In both cases, positive responses meant $Z^s_{i,t+1} = 1$. As an illustration Table 3 shows the correlation, conditional on time fixed effects, household fixed effects and the attrition inverse mills ratio, between prevalent shocks and food security. As an inherently subjective measure, tracking household's perception of shocks allows us to use a household's own internal accounting to approximate how that shock affects a household's well-being. Households could therefore experience multiple shocks at once and fluctuate in and out of experiencing a given shock over time.

Given these two states, experiencing and not experiencing shock s, the probability of passing from state k to state j is a Markov process:

$$Pr\left(Z^s_t = j | Z^s_{t-1} = k\right) = p_{kj} \tag{2}$$

where $k, j \in \{0, 1\}$.

To estimate shock persistence, we use an auto-regressive (AR) linear probability model with one lag.[9]

$$Z^s_{i,t} = \gamma_0 + \gamma^s_1 Z^s_{i,t-1} + \gamma^s_t \left(Z^s_{i,t-1} * \delta_t\right) + \delta_t + \mu^s_i + \epsilon_{i,t} \tag{3}$$

where γ^s_1 conditions the perceived shock s on previously experiencing shock s, γ^s_t allows this persistence to vary by round, δ_t is a monthly time fixed effect and μ^s_i is a household fixed effect. With

[8] It is generally not recommended to use household fixed effects when using a probit maximum likelihood estimator, see Wooldridge (2010).

[9] Estimates of Eq. (3) are robust to 2 and 3 month lags, see supplementary material.

Table 3
Shocks and Food Security.

	Coping Strategy Index			
	(1)	(2)	(3)	(4)
Drought	2.334**			
	(1.027)			
Flood		3.852**		
		(1.672)		
Illness			0.792	
			(1.072)	
Crop Disease				2.772**
				(1.304)
Constant	27.22**	30.83***	28.24**	26.66**
	(10.65)	(10.55)	(10.63)	(10.16)
N	5887	5887	5887	5887

Not reported: time fixed effects δ_t, household fixed effects μ_i^s and inverse mills ratio $\hat{\lambda}$. Village clustered standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

a linear probability model, the coefficients have an intuitive interpretation: $p^s_{0,1} = \gamma^s_0 + \delta_t$ is the probability of experiencing the adverse effects of a given shock, conditional on not experiencing them previously. $p_{1,1} = \gamma^s_0 + \gamma^s_1 + \gamma^s_t + \delta_t$ is the probability that a shock's adverse effects will persist into the next period. γ^s_t and δ_t allow for a non-stationary process since the transition probability can change over time.[10]

We can present our estimated coefficients as a Markov Transition matrix:

	$Z^s_{i,t} = 0$	$Z^s_{i,t} = 1$
$Z^s_{i,t-1} = 0$	$p^s_{0,0} = 1 - (\gamma^s_0 + \delta_t)$ (Probability of not experiencing new shock *s*, given shock *s* was not experienced previously)	$p^s_{0,1} = \gamma^s_0 + \delta_t$ (Probability of experiencing new shock *s*, given shock *s* was not experienced previously)
$Z^s_{i,t-1} = 1$	$p^s_{1,0} = 1 - (\gamma^s_0 + \gamma^s_1 + \gamma'^s_t + \delta_t)$ (Probability of shock *s* not persisting, given shock *s* was experienced previously)	$p^s_{1,1} = \gamma^s_0 + \gamma^s_1 + \gamma'^s_t + \delta_t$ (Probability of shock *s* persisting, given shock *s* was experienced previously)

Our key parameters of interest are $p^s_{1,1}$, the probability of shock *s* persisting, and$p^s_{1,0}$ the probability of recovering from shock *s*. We can think of$p^s_{1,0}$ as resilience to shock *s*. $p^s_{0,1}$ is the probability of shock *s* occurring when it hasn't occurred before. Since it is a subjective measure, we can also think of $p^s_{0,1}$ as vulnerability to shock *s*.

4.1. Estimating shock persistence

We first estimate Eq. (3), the persistence of shocks over time. We take the four most frequent shocks: $S = \{\text{drought}, \text{flooding}, \text{crop disease}, \text{illness}\}$ and regress them against their own values, lagged by one month.[11] Table 4 uses a least squares regression controlling for household and time fixed effects, with errors clustered at the Group Village Head (GVH) level. For succinctness we only report $\hat{\gamma}^s_0$ and $\hat{\gamma}^s_1$ for a given shock *s*. As discussed earlier, since we allow the coefficients to vary over time these reported estimates offer only a snapshot, determined by which time dummy we set as point of reference. Here we set it as May 2017, our last round.[12]

This data also allows us to explore whether multiple shocks interact and whether those effects exacerbate or mitigate one another. Table A3 in the supplementary materials estimates Eq. (3) but includes lags of the other most prevalent shocks. The only significant cross correlation is that experiencing flooding makes it more likely to experience crop disease in the subsequent round.

With four shocks and 11 rounds[13], we can construct a total of 44 transition matrices. For illustrative purposes we choose three periods of reference aligned with the agricultural calendar: November (round 6), when planting begins; February (round 9), the height of the hungry season; and when the harvest comes in May (round 12) (Malcomb, Weaver, & Krakowka, 2014). Table 5 illustrates these transition matrices for drought. Transition matrices for flood, illness and crop-disease are available in the supplementary materials.

Recall that $\hat{p}^s_{1,1} = \hat{\gamma}^s_0 + \hat{\gamma}^s_1 + \hat{\gamma}^s_t + \delta_t$ is the persistence of shock *s*, the probability that a household continues experiencing its adverse effects one month later. As an illustration, the probability of the effects of drought persisting in November is 88.5% (Table 5a). Households that reported drought, and are therefore vulnerable to drought, only had a 11.5% chance of recovering from the effects of drought in the next month. This speaks to how shock's duration and persistence differ. While flooding may only last a week, a drought could last months. Both could have impacts that last months. The persistence of drought does not vary much across seasons, as the probability of a drought's effects persisting are 85.3% in February (Table 5) and 86.1% in May (Table 5c). On the other hand, the persistence of flooding falls and the persistence of crop disease increases as the growing season progresses. Fig. 2a illustrates this change in the persistence of shocks $\hat{p}^s_{1,1}$ visually across all rounds.

In addition to estimating persistence, a useful feature of the specification is that we can separately estimate the probability of a household experiencing the effects of a new shock in terms of its incidence, $\hat{p}^s_{0,1} = \hat{\gamma}^s_0 + \delta_t$. For example, crop-disease starts at a lower level of incidence with $\hat{p}^s_{0,1} = 15.1\%$ in November but this rapidly climbs to 44.8% in February and 41.9% in May. It makes sense that the adverse effects of crop disease are most acute when the planted harvest is nearing maturation. Fig. 2b illustrates how $\hat{p}^s_{0,1}$ changes over time for each shock *s*. We notice that incidence and persistence do not necessarily move in tandem. For example, the effects of drought are very persistent, at $\hat{p}^s_{1,1} > 80\%$ throughout

[10] To avoid collinearity we must set one of the time fixed effects to equal 0. This is arbitrary but then becomes the month of reference for the other δ_t terms.
[11] Out of 15 potential shocks, more than 20% of households reported experiencing these four shocks in at least one month.
[12] This is done by setting the relevant time dummy, i.e. $\delta_{May2017} = 0$ and $\gamma^s_{May2017} = 0$.
[13] We lose one round by construct due to the lag.

Table 4
Lagged Effect of Most Frequent Shocks, OLS.

	(1) Drought	(2) Flood	(3) Illness	(4) Crop Disease
Drought	0.518***			
(1 month lag)	(0.0632)			
Flood		0.125		
(1 month lag)		(0.127)		
Illness			0.405***	
(1 month lag)			(0.0676)	
Crop Disease				0.395***
(1 month lag)				(0.0831)
Constant	0.579*	−0.312	0.0260	0.882**
	(0.278)	(0.331)	(0.254)	(0.401)
N	5036	5036	5036	5036

Not reported: time fixed effects δ_t, interaction of time fixed effects and lagged coefficient γ_t^s, household fixed effects μ_t^s.For reference, we set $\delta_{May2017} = 0$, inverse mills ratio $\hat{\lambda}$. Clustered standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5
Transition Matrices for Drought.

	$Drought_t = 0$	$Drought_t = 1$
(a) *November (Planting Season)*		
$Drought_{t-1} = 0$	$\hat{p}^s_{0,0} = 54.2\%$	$\hat{p}^s_{0,1} = 45.8\%$
$Drought_{t-1} = 1$	$\hat{p}^s_{1,0} = 11.5\%$	$\hat{p}^s_{1,1} = 88.5\%$
Calculated from $\hat{\gamma}^s_0, \hat{\gamma}^s_1, \hat{\gamma}^s_{Nov2016}$ and $\hat{\delta}_{Nov2016}$, where s = drought,in Table 4, column (1)		
(b) *February (Hungry Season)*		
$Drought_{t-1} = 0$	$\hat{p}^s_{0,0} = 58.3\%$	$\hat{p}^s_{0,1} = 41.7\%$
$Drought_{t-1} = 1$	$\hat{p}^s_{1,0} = 14.7\%$	$\hat{p}^s_{1,1} = 85.3\%$
Calculated from $\hat{\gamma}^s_0, \hat{\gamma}^s_1, \hat{\gamma}^s_{Feb2017}$ and $\hat{\delta}_{Feb2017}$, where s = drought,in Table 4, column (1)		
(c) *May (Harvest Season)*		
$Drought_{t-1} = 0$	$\hat{p}^s_{0,0} = 56\%$	$\hat{p}^s_{0,1} = 34\%$
$Drought_{t-1} = 1$	$\hat{p}^s_{1,0} = 13.9\%$	$\hat{p}^s_{1,1} = 86.1\%$
Calculated from $\hat{\gamma}^s_0, \hat{\gamma}^s_1, \hat{\gamma}^s_{May2017}$ and $\hat{\delta}_{May2017}$, where s = drought,in Table 4, column (1)		

the year, where $s = drought$. By contrast $\hat{p}^s_{0,1}$ fluctuates over time, reflecting the seasonality of a drought's incidence.

4.2. Shock persistence and household characteristics

Having mapped out the persistence of various shocks over time, what are the household characteristics correlated with household level shock persistence? From the above specification we can predict the probability of persistence of shock *s* for a given household *i* at time *t*, $\hat{\rho}^s_{i,t} = \hat{p}^s_{0,1} + \hat{\mu}^s_i$. We can then regress this predicted value against time invariant household co-variates X_i outlined in Table 2:

$$\hat{\rho}^s_{i,t} = \alpha_0 + \alpha_1 X_i + \zeta_{i,t} \quad (4)$$

This regression allows us to infer the correlation between these fixed effects and household characteristics $\frac{\delta \hat{\rho}^s_{i,t}}{\delta X_i} = \hat{\alpha}_1$.[14] A negative correlation indicates that households with these characteristics are less likely to experience the persistent effects of a given shock. As discussed earlier, we consider such non-persistence of a shock's effects as a sign of resilience. We estimate Eq. (4) using the observed characteristics from the baseline. Fig. 3 presents estimated correlations with their bootstrapped standard errors. These descriptives are inherently endogenous, but can help inform which type of households are more vulnerable to shocks.

[14] Since these are point estimates, we bootstrap the above two-step process.

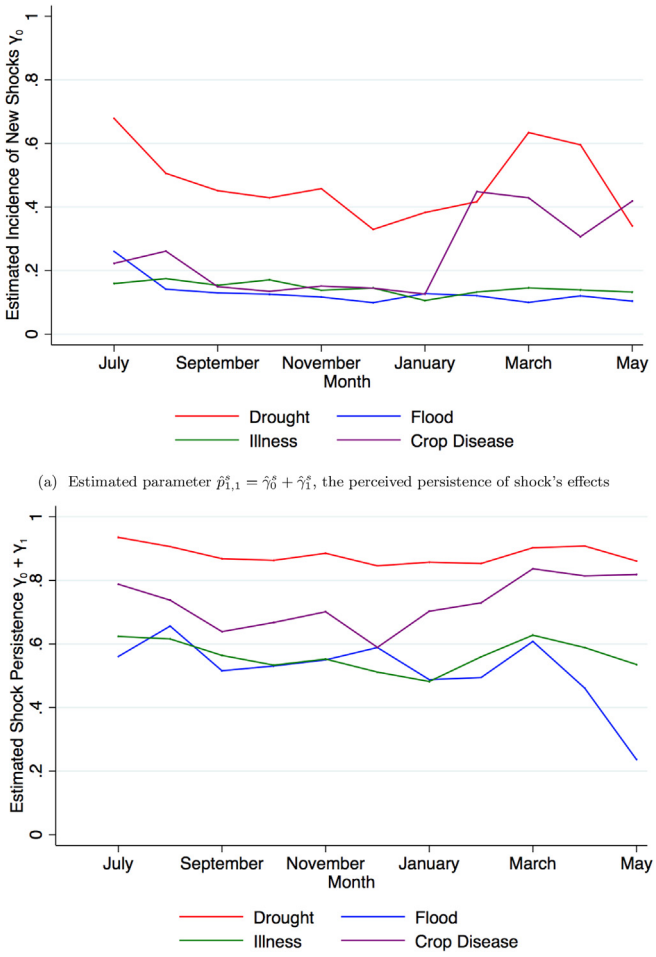

(a) Estimated parameter $\hat{p}^s_{1,1} = \hat{\gamma}^s_0 + \hat{\gamma}^s_1$, the perceived persistence of shock's effects

(b) Estimated parameter $\hat{p}^s_{0,1} = \hat{\gamma}^s_0$, the perceived incidence of new shocks

Fig. 2. Illustration of the parameters calculated from Table 4 as they vary by round.

Living in a flood plain is negatively correlated with the persistence of a drought's adverse effects, but positively correlated with the persistence of illness's adverse effects. Both of these make sense, as soil in the flood-plain is likely to retain more moisture, but the abundance of stagnant water offers breeding pools for malaria and cholera. Having a secondary house often indicates that the household has fields far from its primary home, where they sometimes have to spend the night. This suggests that households

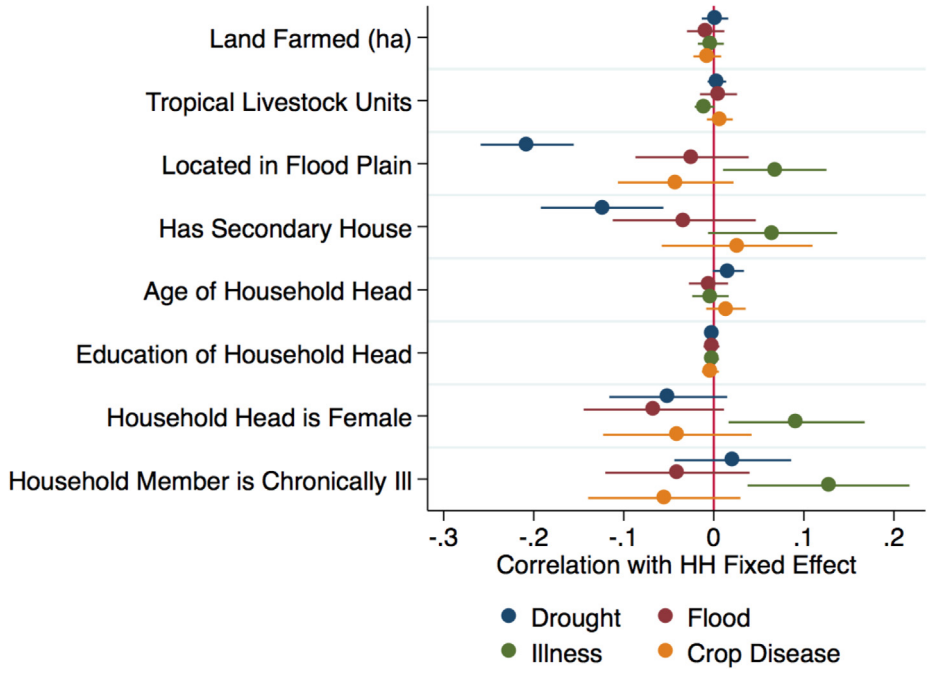

Fig. 3. Correlation between estimated shock specific probability $\hat{\rho}^{s}_{i,t}$ and household characteristics (bootstrapped s.e).

with spatially dispersed fields are more resilient to drought. We find little evidence that the perceived incidence or persistence of a shock is correlated with a household's wealth as measured both in terms of access to land and livestock. Female headed households are more likely to experience the persistent effects of illness and are therefore less resilient. Finally, we find a significant correlation between the persistence of an illness's effects and whether the household has a chronically ill member.

5. Resilience as a moments-based approach

An alternative to looking at the persistence of specific shocks is to think of resilience as the stochastic distribution in food insecurity outcomes $y_{n,t}$ using the Cisse and Barrett (2018) moments based approach. While the above approach emphasizes the persistence of subjective shocks, the moments based approach focuses on the stochastic distribution of well-being expressed as a probability. Here well-being is measured as food insecurity using CSI. Under the first approach, households are resilient to a given shock if the subjective shock experienced is non-persistent. Under the moments based approach, resilient households are those that maintain their well-being above a certain threshold. One focuses on the experience of specific shocks, while the other emphasizes well-being that can be influenced by multiple shocks and household characteristics.

Constructing resilience as a probability using the Cisse and Barrett (2018) method is a three step process:

The first step estimates

$$CSI_{it} = \sum_{\gamma=1}^{3}\left(\beta_{M,\gamma} CSI^{\gamma}_{i,t-1}\right) + \delta_{M,1} X_{i,t} + \delta_{M,2} Z_{i,t-1} + u_{i,t} \tag{5}$$

where $Y_{i,t}$ is CSI, $X_{i,t}$ the set of household co-variates from Table 2 and $Z_{i,t}$ the set of shocks experienced.

The second step regresses the same specification on estimated variance, $\sigma^2_{CSI} = \hat{u}^2_{i,t}$:

$$\sigma^2_{CSI} = \sum_{\gamma=1}^{3}\left(\beta_{V,\gamma} CSI^{\gamma}_{i,t-1}\right) + \delta_V X_{i,t} + \epsilon_{i,t} \tag{6}$$

We posit $\hat{CSI}_{it}$ and $\hat{\sigma}^2_{i,t}$ as the first and second moments of a conditional distribution, respectively. Under the assumption that $CSI_{i,t}$ has either a normal or gamma distribution, we construct the cumulative density function and calculate $\hat{p}_{i,t}(X,Z) = P(CSI_{i,t} \leqslant \bar{CSI})$, where $\bar{CSI}$ is a threshold level of food insecurity. Because CSI is a level of food insecurity, we want to calculate household's probability of being below a given threshold. We posit $\bar{CSI} = 10$. The estimated parameters are reported in Table 6a.[15] We can then estimate

$$\hat{p}_{i,t} = \sum_{\gamma=1}^{3}\left(\beta_{R,\gamma} CSI^{\gamma}_{i,t-1}\right) + \delta_R X_{i,t} + \eta_{i,t} \tag{7}$$

Where $\hat{p}_{i,t}$ can be thought of as 'resilience', the probability that household *i*'s food insecurity will be below $\bar{CSI} = 10$ at time *t*. From Table 6a, the average probability within the sample is 14% if we assume a normal distribution and 13% if we assume a gamma distribution.

5.1. Estimating a moments-based approach

Recalling Cisse and Barrett (2018), estimates of Eqs. (5)–(7) for both a normal and gamma distribution are reported in Table 6b. Of the household covariates, flood plain has a large and significant positive correlation with household resilience. The correlation with having a secondary house is large in magnitude and positive but not statistically significant. The other characteristics we identified as important relative to specific shocks, such as gender and being chronically ill, are not significantly different from 0. Of the shocks, experiencing flooding is correlated with a significant decrease in resilience, as does drought, though the effect is only significant for $\widehat{P}_{Gamma}$. Illness is marginally significant and negative.

These results highlight the similarities and differences between an approach focusing on subjective shock persistence and one based on resilience as a stochastic distribution of food insecurity outcomes. The insights are broadly similar: they emphasize the

[15] Results were robust to alternate weights for CSI and alternative thresholds, see supplementary materials.

Table 6
Estimating Resilience using a Conditional Moments Approach.

(a) Estimated Resilience Parameters

variable	mean	standard deviation	min	max
$\hat{CSI}$	25	8.4	5.9	75
$\hat{\sigma}^2_{CSI}$	121	30	48	252
$\widehat{P}_{Normal}$	0.14	0.16	6.9e-06	0.71
$\widehat{P}_{Gamma}$	0.13	0.2	2.4e-14	.81

We posit $\bar{CSI} = 10$ to calculate $\widehat{P}_{Normal}$ and $\widehat{P}_{Gamma}$.

(b) Resilience, Household Covariates and Reported Shocks

	(1) $\hat{CSI}$	(2) $\hat{\sigma}^2_{CSI}$	(3) $\widehat{P}_{Normal}$ P(CSI < 10)	(4) $\widehat{P}_{Gamma}$ P(CSI < 10)
CSI Last Month	0.093***	0.026***	−0.161***	−0.166***
	(0.002)	(0.001)	(0.024)	(0.030)
Land Farmed (HA)	−0.001	−0.019***	−0.010	−0.011
	(0.002)	(0.001)	(0.026)	(0.030)
Tropical Livestock Units	−0.005***	−0.008***	0.011	0.015
	(0.001)	(0.001)	(0.019)	(0.021)
Flood Plain	−0.062***	0.004	0.235**	0.326***
	(0.006)	(0.003)	(0.107)	(0.124)
Secondary House	−0.011	0.170***	0.160	0.211
or Shack	(0.008)	(0.003)	(0.120)	(0.134)
Age of HH Head	0.000**	0.002***	0.001	0.001
	(0.000)	(0.000)	(0.003)	(0.004)
Education of HH Head	0.001	−0.000	−0.002	−0.002
	(0.001)	(0.000)	(0.013)	(0.015)
HH Head is Female	−0.032***	−0.053***	0.063	0.093
	(0.007)	(0.003)	(0.117)	(0.131)
Chronically Ill	−0.005	−0.056***	−0.039	−0.048
HH Member	(0.008)	(0.004)	(0.129)	(0.145)
Drought	0.129***	0.350***	−0.169	−0.234*
	(0.008)	(0.004)	(0.109)	(0.120)
Flood	0.067***	−0.165***	−0.467***	−0.606***
	(0.007)	(0.003)	(0.139)	(0.166)
Illness	0.060***	0.043***	−0.159	−0.231*
	(0.007)	(0.003)	(0.113)	(0.128)
Crop Disease	−0.026***	−0.061***	0.037	0.049
	(0.006)	(0.003)	(0.101)	(0.115)
N	5011	5011	5011	5011

Not reported: CSI^2, CSI^3, inverse mills ratio $\hat{\lambda}$ Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

persistent adverse effects of flooding and drought on household food security. Both also highlight that living in the flood plain and having a secondary house mitigates these adverse effects. Some of the details differ: Gender and being chronically ill have no significant effect on $\hat{p}_{i,t}$ in Table 6. There are two possible explanations: those characteristics only affect the persistence of illness for the 20–25% of households experiencing illness at any given time, so the effect might be lost in the statistical noise. Alternatively, since persistent illness has little to no effect on food security (see Table 3), if gender and being chronically ill only affects illness, they may not affect the stochastic distribution of food security. This complementary analysis showcases how the two approaches inform our understanding of the realities on the ground, offering a fuller picture of household resilience.

6. Predicting food insecurity

An alternative approach to measuring resilience is using proxies which predict food insecurity in the immediate future. In the reduced form specification above we based our choice of observable characteristics on informed priors, which is not necessarily desirable when we seek to maximize predictive accuracy. Instead we propose to search through the full scope of available data to find the best predictors of future CSI. We do this using supervised machine learning algorithms. This allows us to harness these powerful algorithms to explore the dataset for potentially unexpected predictors, and is to our knowledge the first attempt at using machine learning to inform the resilience literature.

6.1. Prediction using machine learning algorithms

We compare two algorithms, the Least Absolute Shrinkage and Selection Operator (LASSO) and Random Forest. Both allow us to identify the best predictors by selecting a subset of promising variables within a larger set. We also show how the predicted outcomes from these algorithms can be used for geographic targeting.

The use of these algorithms has recently gained popularity in terms of both predicting poverty and shortlisting variables for the purpose of targeting. In terms of targeting, Jean et al. (2016) show that running an image recognition algorithm through publicly available satellite imagery significantly improves geographic poverty targeting.Blumenstock, Cadamuro, and On (2015) use Call Detail Records (CDR) from respondents' cell-phones to predict poverty in Rwanda. A complimentary body of work is concerned with improving accuracy of targeted social programs when lacking comprehensive data on income and consumption. McBride and Nichols (2016) present evidence that applying machine learning algorithms to PMT development can substantially improve the out-of-sample performance of these targeting tools. Kshirsagar et al. (2017) use a bootstrapped LASSO to select a subset of indicators

which accurately predict poverty rates, and show that they outperform a random distribution benchmark. Building on this literature, we innovate here by applying this approach to dynamic food security data.

Though a wide array of popular supervised machine learning tools exist, the general premise is straightforward.[16] First, divide the dataset in two subsets. The first subset, the 'training' set, is used to calibrate the model. The calibrated model is then used to predict outcomes in the second, 'testing' subset of the data. The performance of an algorithm is judged by its predictive accuracy in this second set, as measured using the R^2, weighed against a cost function to penalise the number of covariates. The process is then iterated in an attempt to improve performance. In addition, as an intermediary step some algorithms explicitly identify a subset of variables that are considered the best predictors. These are the predictors we are interested in.

We use the Least Absolute Shrinkage and Selection Operator (LASSO) and Random Forests to identify the best predictors of CSI. LASSO was selected because it retains much of the structure of linear regression analysis, allowing for intuitive interpretation of the coefficients. These can be compared in sign and magnitude to results elsewhere in the paper. Random Forest offers a salient contrast, as by design it is non-linear and every variable is implicitly allowed to interact with every other variable. Though more flexible, this 'black box' approach makes it difficult to interpret the sign of any one coefficient, as it depends on all the others. Comparing the performance of both is therefore useful. More sophisticated deep learning algorithms using neural networks were too demanding of the data given our limited sample size.

We trained our data on the 10 rounds from June to March (the 'training' set) and sought to predict the likely outcomes for April and May (the 'test' set), comparing it with actual CSI levels in those two months.[17] A feature of using this machine learning technique is that it harnesses all the variables collected, rather than just the data described in Section 2. These included asset indicators from the baseline, such as quality of the home, distance to drinking water and diet. They also included time varying indicators, such as the type and source of assistance a household received and any change in assets. We kept all variables with less than 2% of observations missing.[18] To account for attrition, we only kept the households for which we had observations in April and May, accounting for the drop in N in Tables 7 and 8. After this cleaning process, we were left with an $\mathbb{R}^d$ space of predictors, where d = 79 variables.

6.2. LASSO

In a traditional regression, additional parameters always increase predictive performance but risk overfitting the data. LASSO, short for Least Absolute Shrinkage and Selection Operator, is a linear regression which penalizes additional parameters β by including the term $\kappa\|\beta\|$.[19] The modified least squares operator is therefore

$$\min_{\beta\in\mathbb{R}^d}\left\{\frac{1}{N}\sum_{i=1}^{N}\left(Y_{i,t}-\beta X_{i,t}\right)^2+\kappa\|\beta\|\right\} \quad (8)$$

Table 7
Top Variables Selected by LASSO Algorithm.

Variable	Coefficient
CSI Last Month	0.473***
	(0.015)
Group Village Head	0.012***
	(0.002)
Received Assistance from Government Last Month	−1.535**
	(0.727)
Received Assistance as Food Last Month	−0.662*
	(0.417)
Distance to Drinking Water (minutes walking)	0.013*
	(0.007)
Quality of Floor	−0.518
	(0.368)
Sold Assets Last Month	−0.818
	(0.691)
Purchased Assets Last Month	−0.778
	(0.694)
Lives in Flood Plain	−0.662
	(0.562)
Experienced Drought Last Month	0.434
	(0.371)
N	4308
Out of sample R^2 (April, May)	56.4%

Sample restricted to training set, (June–March) Bootstrapped standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8
Top Variables Selected by Random Forest Algorithm.

Variable	ΔMSE^* (standardized)*
CSI Last Month	21.70
Group Village Head	25.58
Age	19.98
Education	17.32
Land Farmed (ha)	17.91
Dietary Diversity Score	17.02
Distance to Drinking Water (minutes walking)	14.27
Quality of Roof	16.76
Pregnant or Nursing Household Member	14.65
Lives in Flood Plain	13.80
N	4308
Out of sample R^2 (April, May)	55.6%

Sample restricted to training set, (June–March).
*Increase in MSE when variable is omitted, measure of variable importance.

Each $\beta \in \mathbb{R}^d$ must therefore add sufficient explanatory power to overcome the penalty term $\kappa\|\beta\|$, otherwise it is minimized to 0. Since we have no good a priori for the penalty term κ the algorithm we use glmnet, which uses coordinate descent and a soft threshold operator to iterate through a plausible range (Friedman, Hastie, & Tibshirani, 2010). For each candidate κ, the algorithm randomly subsets the training data further and computes the mean cross-validated mean squared error (MSE). It settles on two candidates values of $\kappa : \hat{\kappa}^{min}$ minimizes the mean cross-validated error. $\hat{\kappa}^{1se}$ is the most parsimonious model in terms of number of parameters β while within a standard error of the minimum. Since we seek to identify the subset of best predictors, we report the results from $\hat{\kappa}^{1se}$, the parsimonious model.

In order to identify the best predictors of CSI, we bootstrap the training data and run the LASSO algorithm through a thousand

[16] 'Supervised' machine learning use inputs x to predict outputs y. 'Unsupervised' tools seek to identify patterns in x without corresponding outputs.
[17] Best practice is to split the data into roughly 80% training and 20% testing. We erred on the side of a slightly larger training dataset because the algorithms further subdivide this dataset for cross-validation.
[18] In order to avoid dropping too many observations, we substituted values for those missing variables using nearest matches from observed data, randomly sampling from the set of nearest observations.
[19] We use κ to differentiate from the inverse mills ratio λ, defined earlier.

iterations, similar to Kshirsagar et al. (2017). We then compute the mean coefficient and its standard deviation and keep the 10 most significant variables.[20] We report these predictors in Table 7. Lagged CSI is unsurprisingly a strong predictor of future CSI. Location, as indicated by community id (GVH), is also a strong predictor. Receiving assistance from the government and receiving it as food are both predictors of future decreases in CSI, and therefore decreased food insecurity. The algorithm selects four types of asset indicators: distance to drinking water, whether any assets were bought, whether assets were sold, and the quality of the floor. Every additional minute of walking distance to drinking water increases predicted CSI. Interestingly both buying and selling assets are predictors of decreased CSI, suggesting that what matters is liquidity. Living in a flood plain decreases predicted CSI, and experiencing drought last month increases it.

With the exception of living in a flood plain, these predictors do not directly correspond to the characteristics we used in our earlier specification. Unlike those characteristics, five of these predictors vary over time and are therefore good proxies for a household's fluctuating state of food insecurity in the immediate. This points to a difference in time frame. Whilst our earlier approaches sought to estimate characteristics contributing to resilience over a year or half year time frame, this exercise emphasizes predictive accuracy over a one to two month time frame. Time frame therefore affects the characteristics of interest to our resilience analysis.

We then use the estimated parameters to predict CSI in our test data and compare its performance to the actual CSI outcomes. When compared, our LASSO algorithm gives us an $R^2 = 56.4\%$.

6.3. Random forest

While it does allow us to narrow down the set of predictors, a LASSO algorithm remains a linear regression. An alternative approach, regression trees, offers a more flexible functional form (Hastie, Tibshirani, & Friedman, 2009). A regression tree chooses variables in the set $x_j \in X$ and the value of that variable s, that split the set into two 'branches' which form half planes R_1 and R_2. x_j and s are chosen through an optimization process which jointly minimizes the mean squared error for the dependent variable y_i in each of the half planes defined by the branches.

$$\min_{js}\left[\min_{c_1}\sum_{x_i\in R_1(j,s)}(y_i-c_1)^2+\min_{c_2}\sum_{x_i\in R_2(j,s)}(y_i-c_2)^2\right] \tag{9}$$

where

$$c_1=\frac{1}{n}\sum_i(y_i|x_i\in R_1(j,s)) \quad \text{and} \quad c_2=\frac{1}{n}\sum_i(y_i|x_i\in R_2(j,s)) \tag{10}$$

Each half plane is in turn split into two branches and the process is iterated to create a 'tree' which fits the data. Single regression trees tend to suffer from over-fitting, mistakenly attributing random variations in the outcome to an explanatory variable. In order to correct for this, random forest randomly select a subset of$x_j \in X$ as candidate variables for the regression tree, which it fits to the data (Breiman, 2001). By repeating the process, it can compare the cross-mean performance across these 'trees', hence creating random forest. Because the performance of any given variable is conditional on its parent branches, regression trees do not allow for explicit coefficient estimates. Instead variable x_j's performance is measured as ΔMSE_j, ie the increase in Mean Squared Error if the variable is omitted.

We bootstrap the training data and run the random forest algorithm through a thousand iterations. This gives us a thousand values for each variable: ΔMSE_j^b. We calculate the mean across bootstrapped values and divide by the standard deviation to normalize it. This allows us to rank the 10 best predictors of CSI, listed in Table 8. Note that because these are not coefficients, we cannot say whether individual variables predict increased or decreased CSI. Like LASSO, it finds that CSI last month, location and the distance to drinking water are good predictors. Indeed five of the top ten variables overlap between the two approaches. Unlike LASSO, the regression tree algorithm favors household characteristics from the baseline. Many of these correspond to the characteristics we used in our other specifications, including age, education, land farmed and whether the household lived in a flood plain. A household's dietary diversity score is also a good predictor.[21] The quality of a households roof and whether anyone in the household is pregnant or nursing round of the top ten variables selected.

By running the subset of selected variables through one more iteration of a random forest, we generated out of sample predictions for April and May, giving us an $R^2 = 55.6\%$.

Fig. 4 compares the predictors across both algorithms. Five of the top ten variables overlap or are very similar. Previous month's CSI is the best predictor in both cases, followed closely by geographic location as indicated by group village head. Other good predictors include the household's distance to drinking water, the quality of their home and whether they live in a floodplain. The differences between what each algorithm picks up are illustrative as well, and speaks to the differences in the objective function optimized. LASSO, by weighing each variable as a stand-alone in a linear regression, favors time varying variables that immediately affect CSI. By optimizing iteratively, regression trees implicitly condition variables on each other, thereby favoring underlying variables which affect CSI via their effect on other variables. Hence it favors time invariant characteristics. In terms of out-of-sample performance, the two methods are comparable with an R^2 of 56.4% for LASSO and 55.6% for Random Forest. The predicted values account for more than half of the observed variation in food security outcomes across households. Though illustrative of the predictive potential of combining high-frequency data with machine-learning algorithms, there is room to improve this performance by expanding the set of variables to include geo-spatial indicators, an area of future research.

6.4. Mapping predicted outcomes

Since this data collection exercise occurred during a humanitarian emergency, we worked with our partners on the ground to feed data into their decision making process. Given the uneven pace of households recovery from drought, we sought to determine if there were lingering 'pockets' of food insecurity. We used the predictions from the above algorithms to map out the predicted CSI levels $\hat{CSI}$, available in the supplementary materials. The figure maps the actual CSI outcomes in April and May, and compares it to the predicted outcomes using the LASSO and Random Forest algorithms. This allowed us to demonstrate the usefulness of a high resolution system of embedded enumerators.

This approach highlights how decision makers could target interventions with improved accuracy. Though not always precise in magnitude, the predictions provide an accurate forecast of where we could expect high levels of CSI. These tended to be in tightly circumscribed geographical areas. Such concentrated levels of high CSI were due to localized co-variate shocks: some commu-

[20] We kept non-statistically significant variables for their potential predictive power.

[21] Dietary diversity score, or DDS, is the sum of food types a household reports having consumed in the past 24 h.

	LASSO	Random Forest
	CSI Last Month	CSI Last Month
	Group Village Head	Group Village Head
	Received Assistance from Government Last Month	Age
	Received Assistance as Food Last Month	Education
	Distance to Drinking Water	Distance to Drinking Water
	Quality of Floor	Quality of Roof
	Sold Assets Last Month	Dietary Diversity Score
	Purchased Assets Last Month	Land Farmed (ha)
	Lives in Flood Plain	Lives in Flood Plain
	Experienced Drought Last Month	Pregnant or Nursing Household Member
N	4308	4308
R^2	56.4 %	55.6 %

Sample restricted to training set, (June-March) and tested on April-May

Fig. 4. Top predictors of CSI.

nities were still struggling to recover from the effects of drought. Communities in the flood plain experienced the adverse effects of localized flooding, while other communities experienced outbreaks of crop disease.

7. Discussion

The present study was designed to examine how selected approaches to resilience measurement can generate different sets of insights about the dynamic relationship between experienced shocks and well-being. The approaches we tested reflect different strategies with reference to shocks, stochastic probabilities, and the prediction of future well-being. Across these approaches, a comparison of findings demonstrates both the unique offering of each method and their complementary insights.

The paper starts with an analysis estimating the persistence of subjective shocks, modelled as a non-stationary Markov Matrix. This allows a distinction between shock incidence and shock persistence, which do not necessarily move in tandem. In being subjective, the incidence and persistence of these shocks reflects their effects on household well-being. Some shocks, like drought, are very persistent in their effects, while the persistence of shocks like flood and crop disease vary over time. It further tests whether observed household characteristics are correlated with the estimated persistence of specific shocks. It finds that having secondary sheds, which was used as a proxy indicator for the spatial dispersion of fields, is correlated with a lower persistence of drought's adverse effects. Living in the flood plain is also correlated with lower persistence of drought's adverse effects. Female headed households and households with a chronically ill member experience more persistent effects of illness.

To consider the benefits derived from the comparative approach employed in the present study, it is useful to consider the goals of resilience measurement. If the goal is to unpack shock-specific dynamics, the first method based around shocks is ideal for an exploratory analysis linking household characteristics to their resilience to specified shocks. This can in turn inform program design. Only some of these characteristics, however, are relevant if, instead of subjective shock persistence, one is interested in looking at the effects of household characteristics and shocks on household food insecurity. Here, the Cisse and Barrett (2018) moments-based approach may be the best option. Consistent with findings from the persistence analysis, which finds that living in a flood plain is related to improved resilience, the moments-based approach finds that living in the flood plains significantly increases the probability that a household remains below a threshold of food insecurity. The approach does not however identify secondary sheds or female headed households as important drivers of resilience, perhaps because these are specific to certain idiosyncratic shocks rather than affecting well-being overall. This demonstrates how the approaches differ but complement one another. By providing a more holistic metric of resilience, the moments based approach is well-suited for impact evaluation, with the intervention on the right hand side of the equation to determine whether it had a significant effect on resilience.

Following the third approach applied in the present study, the data can be used to identify the best predictors of food insecurity in the immediate future, in months following a shock. Rather than rely on a subset of selected explanatory variables, these algorithms explore the full scope of the dataset for predictors. Since these predictors mediate the effect of shocks on food insecurity, we consider them proxies of resilience. To our knowledge such an approach has not yet featured in the resilience literature. Two different

algorithms are harnessed for the task, the Least Absolute Shrinkage and Selection Operator (LASSO) and Random Forest algorithm. The best predictors across both algorithms are previous levels of food insecurity, living in a flood plain and distance to drinking water. This list of top predictors can be used to inform future low-burden, high-frequency data collection exercises. Mapping out these predictions and comparing them to actual outcomes, we find, with high accuracy that high levels of food insecurity are concentrated in small pockets. This finding reflects the local nature of most subjective shocks. This can inform program targeting based on geography or observable household criteria.

In conclusion, we believe that the insights gained from the present study demonstrate the value of multiple complementary approaches for resilience measurement. All three approaches consistently emphasized the persistence of shocks and food insecurity, suggesting that households struggle to recover. They also highlighted the importance of living in the flood plain as a driver of resilience. The results obtained from the approaches, however, differed in potentially relevant ways for the purpose of designing interventions. If drought is the major concern, secondary sheds and the geographic dispersion of fields they imply matter. If illness is a concern, then the gender of the household head is important. If we care about forecasting food insecurity, proxies such as distance to drinking water should be accounted for. Finally, as an added methodological contribution, we have sought to demonstrate the value-added of monthly high-frequency data. We hope to leverage the existing network of embedded enumerators to create a system of sentinel sites, representative communities whose outcomes in terms of shocks and food-security offer real-time snapshots of the situation on the ground. These features of resilience-focused investigations can both improve the accuracy of early warning systems in humanitarian emergencies and strengthen the quality of resilience measurement.

Acknowledgements

Special thanks to John Hoddinott, Chris Barrett and Participants at the AAEA 2018 Annual Conference and World Bank Brown Bag Seminar for their feedback on earlier versions of this paper.

This work was funded with support from Catholic Relief Services. Thank you to James Campbell and the implementation team at Catholic Relief Services for their invaluable assistance in piloting and implementing the research protocol.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, athttps://doi.org/10.1016/j.worlddev.2019.04.010.

References

Alinovi, L., Mane, E. & Romano, D. (2009). Measuring household resilience to food insecurity: Application to palestinian households. *EC-FAO Food Security Programme: Linking Information and Decision Making to Improve Food*.

Angeon, V., & Bates, S. (2015). Reviewing composite vulnerability and resilience indexes: A sustainable approach and application. *World Development, 72*, 140–162.

Arouri, M., Nguyen, C., & Youssef, A. B. (2015). Natural disasters, household welfare, and resilience: Evidence from rural vietnam. *World Development, 70*, 59–77.

Bank, W. (2010). World development report 2010 : Development and climate change. World Bank. Retrieved from https://openknowledge.worldbank.org/handle/10986/4387.

Barrett, C., & Santos, P. (2014). The impact of changing rainfall variability on resource-dependent wealth dynamics. *Ecological Economics, 105*, 48–54.

Barrett, C. B., & Constas, M. A. (2014). Toward a theory of resilience for international development applications. *Proceedings of the National Academy of Sciences, 111* (40), 14625–14630.

Barrett, C., Smith, K., & Box, P. (2001). Not necessarily in the same boat: Heterogeneous risk assessment among east african pastoralists. *Journal of Development Studies, 37*(5), 1–30.

Bé né, C., Wood, R., Newsham, A. & Davies, M. (2012). Resilience: New utopia or new tyranny? reflection about the potentials and limits of the concept of resilience in relation to vulnerability reduction programmes. *IDS Working Papers*, 2012 (405), 1–61.

Bé né, C., Chowdhury, F. S., Rashid, M., Dhali, S. A., & Jahan, F. (2017). Squaring the circle: Reconciling the need for rigor with the reality on the ground in resilience impact assessment. *World Development, 97*, 212–231.

Blumenstock, J., Cadamuro, G., & On, R. (2015). Predicting poverty and wealth from mobile phone metadata. *Science, 350*(6264), 1073–1076.

Breiman, L. (2001). Random forests. *Machine Learning, 45*(1), 5–32.

Brown, C., Ravallion, M., & Van De Walle, D. (2018). A poor means test? Econometric targeting in africa. *Journal of Development Economics, 134*, 109–124.

Brown, M., Funk, C., Galu, G., & Choularton, R. (2007). Earlier famine warning possible using remote sensing and models. *EOS, 88*(39).

Brück, T., d'Errico, M., & Pietrelli, R. (2018). The effects of violent conflict on household resilience and food security: Evidence from the 2014 gaza conflict. *World Development*.

Carter, M., & Barrett, C. (2006). The economics of poverty traps and persistent poverty: An assetbased approach. *Journal of Development Studies, 42*(2), 178–199.

Carter, M., Little, P., Mogues, T., & Negatu, W. (2007). Poverty traps and natural disasters in ethiopia and honduras. *World development, 35*(5), 835–856.

Chavas, J.-P. (2017). Agro-ecosystem productivity and the dynamic response to shocks. *The economics of poverty traps*. University of Chicago Press.

Cisse, J. D., & Barrett, C. B. (2018). Estimating development resilience: A conditional moments-based approach. *Journal of Development Economics*.

Constas, M., Frankenberger, T., & Hoddinott, J. (2014a). *Resilience measurement principles, food security information network technical series (tech. rep. No. 1)*. Rome: Food, Agriculture Organization and World Food Programme.

Constas, M., Frankenberger, T., Hoddinott, J., Mock, N., Romano, D., Bene, C., & Maxwell, D. (2014b). *A proposed common analytical model for resilience measurement: A general causal framework and some methodological options (tech. rep. No. 2)*. Rome: Food, Agriculture Organization and World Food Programme.

Dercon, S. (2006). Vulnerability: A micro perspective. *Securing Development in an Unstable World, 30*, 117–146.

d'Errico, M., & Di Giuseppe, S. (2018). Resilience mobility in uganda: A dynamic analysis. *World Development, 104*, 78–96.

d'Errico, M., Grazioli, F., & Pietrelli, R. (2018). Cross-country evidence of the relationship between resilience and the subjective perception of well-being and social inclusion: Evidence from the regions of matam (senegal) and the triangle of hope (mauritania). *Journal of International Development, 30*(8), 1339–1368.

d'Errico, M., Romano, D., & Pietrelli, R. (2018). Household resilience to food insecurity: Evidence from tanzania and uganda. *Food Security, 10*(4), 1033–1054.

Doss, C., McPeak, J., & Barrett, C. B. (2008). Interpersonal, intertemporal and spatial variation in risk perceptions: Evidence from east africa. *World Development, 36* (8), 1453–1468.

Forsyth, T. (2018). Is resilience to climate change socially inclusive? Investigating theories of change processes in myanmar. *World Development, 111*, 13–26.

Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software, 33*(1), 1.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed). New York: Springer.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica, 47* (1), 153.

Hoddinott, J., & Kinsey, B. (2001). Child growth in the time of drought. *Oxford Bulletin of Economics and Statistics, 63*(4), 409–436.

Holling, C. S. (1973). Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics, 4*(1), 1–23.

Hunter, B., Gray, M., & Edwards, B. (2013). The use of social surveys to measure drought and the impact of drought. *Social Indicators Research, 113*(1), 419–432.

IPCC (2018). *An ipcc special report on the impacts of global warming of 1.5c above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty*. In: V. Masson-Delmotte, P. Zhai, H. O. Pörtner, D. Roberts, J. Skea, P. R. Shukla,... W. T. (eds.). World Meteorological Organization.

Jean, N., Burke, M., Xie, M., Davis, W. M., Lobell, D. B., & Ermon, S. (2016). Combining satellite imagery and machine learning to predict poverty. *Science, 353*(6301), 790–794.

Kshirsagar, V., Wieczorek, J., Ramanathan, S. & Wells, R. (2017). Household poverty classification in data-scarce environments: A machine learning approach. arXiv preprint arXiv:1711.06813.

Kuhl, L. (2018). Potential contributions of market-systems development initiatives for building climate resilience. *World Development, 108*, 131–144.

Le Houérou, H., & Hoste, C. (1977). Rangeland production and annual rainfall relations in the mediterranean basin and in the african sahelo-sudanian zone. *Journal of Range Management*, 181–189.

Lee, J. H., Nam, S. K., Kim, A., Kim, B., Lee, M. Y., Lee, S. M., et al. (2013). Resilience: A metaanalytic approach. *Journal of Counseling & Development, 91*(3), 269–279.

Lillard, L. A., & Panis, C. W. (1998). Panel attrition from the panel study of income dynamics: Household income, marital status, and mortality. *Journal of Human Resources*, 437–457.

Malcomb, D. W., Weaver, E. A., & Krakowka, A. R. (2014). Vulnerability modeling for sub-saharan africa: An operationalized approach in malawi. *Applied Geography, 48*, 17–30.

Matin, N., Forrester, J., & Ensor, J. (2018). What is equitable resilience? *World Development, 109*, 197–205.

Maxwell, D. G. (1996). Measuring food insecurity: The frequency and severity of "coping strategies". *Food Policy, 21*(3), 291–303.

Maxwell, D., Caldwell, R., & Langworthy, M. (2008). Measuring food insecurity: Can an indicator based on localized coping behaviors be used to compare across contexts? *Food Policy, 33*(6), 533–540.

McBride, L., & Nichols, A. (2016). *Retooling poverty targeting using out-of-sample validation and machine learning*. The: World Bank.

Oswald, A. J., & Wu, S. (2010). Objective confirmation of subjective measures of human well-being: Evidence from the usa. *Science, 327*(5965), 576–579.

Quandt, A. (2018). Measuring livelihood resilience: The household livelihood resilience approach (hlra). *World Development, 107*, 253–263.

Ramanath, R. (2016). Defying ngo-ization?: Lessons in livelihood resilience observed among involuntarily displaced women in mumbai, india. *World Development, 84*, 1–17.

Shah, S. H., Angeles, L. C., & Harris, L. M. (2017). Worlding the intangibility of resilience: The case of rice farmers and water-related risk in the philippines. *World Development, 98*, 400–412.

Smerlak, M. & Vaitla, B. (2016). A non-equilibrium formulation of food security resilience. Social Science Research Network. Retrieved from https://ssrn.com/abstract=2792449.

Smith, L. C., & Frankenberger, T. R. (2018). Does resilience capacity reduce the negative impact of shocks on household food security? Evidence from the 2014 floods in northern bangladesh. *World Development, 102*, 358–376.

Stevenson, B., & Wolfers, J. (2013). Subjective well-being and income: Is there any evidence of satiation? *The American Economic Review, 103*(3), 598–604.

Taylor, S. E. (1991). Asymmetrical effects of positive and negative events: The mobilization-minimization hypothesis. *Psychological bulletin, 110*(1), 67.

Thulstrup, A. W. (2015). Livelihood resilience and adaptive capacity: Tracing changes in household access to capital in central vietnam. *World Development, 74*, 352–362.

Tierney, K., & Bruneau, M. (2007). Conceptualizing and measuring resilience: A key to disaster loss reduction. *TR, news*, 250.

Ververs, M.-T. (2012). The east african food crisis: Did regional early warning systems function? *The Journal of Nutrition, 142*(1), 131–133. https://doi.org/10.3945/jn.111.150342. eprint:/oup/backfile/ content_public/journal/jn/142/1/10.3945_jn.111.150342/4/131.pdf.

Walsh-Dilley, M., Wolford, W., & McCarthy, J. (2016). Rights for resilience: Food sovereignty, power, and resilience in development practice. *Ecology and Society, 21*(1).

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.

Zimmerman, F., & Carter, M. (2003). Asset smoothing, consumption smoothing and the reproduction of inequality under risk and subsistence constraints. *Journal of Development Economics, 71*(2), 233–260.